REMARKS

Rejections Under 35 USC 112

In response to the rejections of claims 1, 6 for lack of antecedent basis for "the central process," claim 1 is amended to recited "a central process" rather than "the central process" in step b, hence providing sufficient antecedent basis for the subsequent use of this phrase in claim 1 as well as in claim 6.

With regard to claims 12 and 13, the labeling of the steps in each of these claims is amended to render it consistent with that in claim 1, on which these claims depend.

Further, as listed above, claim 14 is amended to recite "a central process" rather than "the central process" in step b. to overcome the 112 rejection relative to this claim. Further, claim 19 is amended to recite "the central process" in line 10 rather than "the central processor."

The Examiner also rejects claim 23 for lack of antecedent basis for "the central process." It, however, appears that the Examiner intended to raise this rejection relative to claim 24, and not claim 23, as claim 23 does not include a recitation of "the central process." Accordingly, claim 24 is amended, as noted above, to recite "a central process" at line 5.

Rejections Under 35 USC 103

The Office Action rejects claims 1, 2, 3, 4, 8, 12, 13, 15, 16, 17 and 20 as being obvious in view of an article entitled "An Introduction to Novell's Burst Mode Protocol" authored by Stevenson et al. and what the Examiner characterizes as Applicant's admitted prior art.

Claim 1, as amended, recites a method of managing distributed statistical data retrieval in a network device by employing the following steps:

a. gathering statistical data on at least one card within the network device periodically;

b. sending a predetermined number of packets from the cart to a central process, wherein each packet includes at least a portion of the statistical data;

c. sending *an acknowledge request* to the central process in conjunction with sending the last packet in the predetermined number; and

d. controlling the number of packets sent from the card to the central process, including:

sending *an acknowledge packet* from the central process to the card *indicating a time that the card can resume sending packets to the central process*,

repeating steps b, c and d when the acknowledge packet is received at the card

Stevens discloses a protocol, known as Burst Mode Protocol, for transferring data between a client and a server. For example, a client can send a request to a file server to read a file. In response, the file server can send a series of packets (i.e., a burst) to the client. When client receives the last packet, *the burst transaction is complete if there are no missing fragments* (packets). Otherwise, the client sends a list of missing fragments to the server, and the server retransmits these fragments. A number of bursts may be needed to complete the request. A maximum burst window, i.e., the maximum number of packets in a burst, below a theoretical threshold of 128 packets for a particular transmission, can be adjusted based on network traffic congestion conditions. More particularly, the window size can be modified by successes and failures in transmitting packets. For example, a time-out or a dropped packet can be considered a failure and can result in narrowing of the burst window. In addition, the time delay between transmission of packets within a burst can also be adjusted.

In a Burst Mode Protocol, the client does not send an acknowledgement packet to the server, upon receipt of a burst, indicating *a time that the server can resume sending a subsequent burst*. In other words, although Stevenson describes adjusting the time gap between packets within a burst, it does not disclose adjusting the time between the bursts based on a request received from the client. In contrast, claim 1, as amended, recites sending an acknowledge packet from the central process to the card indicating a time that the card can resume sending packets to the central process.

In addition, Stevenson does not teach or even remotely suggest utilizing its protocol for sending *statistical data* from a plurality of *cards* in a network device to a *central process* executing on that device. In fact, Stevenson is primarily directed to protocols for sending packets

between nodes in a network, and not to managing transmission of statistical data from different cards in a network device to a central process executing on that device. In particular, a protocol that may be suitable for inter-device communication within a network may not be necessarily optimal for intra-device communication, e.g., due to protocol complexity designed for dealing with network congestion and topology changes that may not be relevant to communication between cards within a network device and a central process.

Moreover, Applicants take issue with the assertion in the Office Action that certain prior art has been admitted. In the background of the invention, Applicants discuss problems associated with retrieving data from distributed modules within a network device, and in no way indicate any prior art solution for these problems. Applicants note in this regard that it is improper to use hindsight and Applicants own invention as a blue print for modifying prior art to obtain the disclosed invention.

Hence, claim 1 distinguishes patentably over Stevenson. In addition, claim 2 depends on claim 1, and is also patentable. Claim 3 depends on claim 1, and further recites that the step of sending an acknowledge request to the central process in conjunction with sending the last packet in the predetermined number includes sending the acknowledge request in an acknowledge request packet *separate* from the last packet in the predetermined number. In contrast, in the Burst Mode Protocol, an End of Burst flag is included in the last packet of a burst.

Likewise, claims 4, 8 and 12 depend on claim 1, and hence incorporate its patentable features.

Independent method claim 14 is amended in a fashion similar to claim 1 to recite, among other steps, that controlling the number of packets sent from the cards to the central process includes sending an acknowledge packet from the central process to each card indicating *a time at which the card can resume sending packets to the central process* – a feature not taught by Stevenson as discussed in detail above.

Thus, claim 14 and claims 15, 16, 17 and 20 that depend on claim 14 distinguish patentably over Stevenson.

In Paragraph 20, the Office Action rejects claims 5, 9-11, 13, 18, and 21-25 as being obvious over Stevenson in view of what the Examiner again characterizes as Applicant's admissions and further in view of certain teachings of Tananebaum.

Claim 5 is rewritten in an independent format to include the features of claim 1 on which it originally depended. More specifically, claim 5 recites a method of managing distributed statistical data retrieval in a network device the includes the following steps:

a. gathering *statistical data* on at least one *card* within the *network device* periodically;

b. sending a *predetermined number of packets* from the card to a central process, wherein each packet includes at least a portion of the statistical data;

c. *sending an acknowledge request* to the central process in conjunction with sending the last packet in the predetermined number; and

d. *controlling the number of packets* sent from the card to the central process, including:

sending an acknowledge packet from the central process to the card; and

repeating steps b, c and d when the acknowledge packet is received at the card;

wherein sending an acknowledge packet from the central process to the card, comprises: detecting an acknowledge request at the central process in a packet received from the card;

determining a number of packets to be processed by the central process;

comparing the number of packets to be processed to a *predetermined threshold* periodically; and

sending the acknowledge packet to the card from the central process when the number of packets to be processed is less than the predetermined threshold.

As an initial matter, in the Burst Mode Protocol described by Stevenson, the client does not send, upon receipt of a burst, an acknowledge packet to the server *when the number of packets to be processed by the client is less than a predetermined threshold.* Further, the passage in Tanenbaum to which the Examiner refers simply describes a so-called stop-and-wait protocol in which a sender sends a frame to a receiver and then waits for an acknowledgement,

in the form of a dummy frame, from the receiver before transmitting the next frame. The cited passage of Tanenbaum, however, does not teach or suggest determining, upon receipt of a frame by the receiver, a number of packets to be processed by the receiver, comparing that number with a predetermined threshold, and sending the acknowledgement to the sender when the number of packets to be processed is less than the predetermined threshold. In other words, although the cited passage of Tanenbaum describes sending an acknowledgement to the sender, it does not disclose the steps recited in claim 5 for *determining when the acknowledgement should be sent.*

Moreover, both Stevenson and Tanenbaum are primarily directed to protocols for *sending packets between nodes in a network*, and not to managing transmission of *statistical data* from different *cards* in a network device to a *central process* executing on that device. As noted above, a protocol that may be suitable for inter-device communication within a network may not be necessarily optimal for intra-device communication.

Thus, claim 5 is patentable over the combined teachings of the cited art.

Claims 9-11, 13 depend either directly or indirectly on claim 1, and hence incorporate its features. As discussed above, Stevenson fails to teach the salient features of claim 1. Tanenbaum does not bridge the gap in Stevenson. In particular, in the protocol discussed in the cited passage of Tanenbaum, an acknowledgement sent by the receiver to the sender does *not* include a *time* at which the sender can resume transmitting packets to the receiver. In contrast, claim 1 (and hence claims 9-11) recites sending an acknowledge packet from the central process to the card indicating a time that the card can resume sending packets to the central process. Similar arguments apply to establish that claims 18 and 20, which depend on claim 14, also distinguish over the combined teachings of the cited references.

Claim 21 depends on claim 20, and further recites that gathering statistical data on a plurality of cards within the network device periodically comprises adding the current statistical data sample to a data summary on each card each time the current statistical data sample is gathered – a feature not taught by any of the cited references. And claim 22 depends on claim 21 and further recites that sending a predetermined number of packets from each card to the

central process comprises sending packets containing at least a portion of the current statistical data sample from each card to the central process periodically in a first period and sending packets containing at least a portion of the data summary from each card to the central process periodically at a second period. As noted above, Stevenson is not concerned with transmission of statistical data from a plurality of cards within a network device to a central process, much less sending two sets of packets, one set containing at least a portion of the data and another set containing a portion of data summary, during different temporal periods. Hence, claim 22 and claim 23 that depends on claim 22 are also patentable over the cited art.

Independent claim 24 recites a method of managing distributed statistical data retrieval in a network device that includes the steps of gathering plurality of different types of statistical data on at least one card within the network device periodically, and sending groups of packets from the card to the central process at *staggered* times, wherein each group of packets includes one of the different types of statistical data.

Steventson, as noted above, is not directed to managing transmission of statistical data from a card in a network device to a central process executing on that device. Moreover, its protocol does not call for sending data packets at *staggered* times, as recited in claim 24. In addition, there is no indication in Applicant's background regarding suggestions in the art related to such an approach. Hence, claims 24 and 25, which depends on claim 24, distinguish over the cited art.

In Paragraph 28, claims 6 and 19 are rejected as being unpatentable over Stevenson in view what the Examiner characterizes as Applicant's admission and further in view of U.S. Patent No. 6,167,054 of Simmons.

Claim 6 depends on claim 1 and further recites that sending an acknowledge packet from the central process to the card further includes detecting an acknowledge request at the central process in a packet received from the card, determining a number of packets to be processed by the central process, estimating when the number of packets to be processed will be below the determined threshold.

As discussed in detail above, Stevenson's protocol does not require the client to send an acknowledge request to the server to indicate a time at which the server can resume transmitting packets, and further does not require the client to determine the number of packets to be processed and to estimate when the number of packets will be below a determined threshold for instructing the server to continue transmission.

Simmons describes a flow control method for managing transmission of data frames from a transmitting station to a destination station via one or more switches. In Simmon's method, a switch can transmit a PAUSE signal to the transmitting station to stop transmission for a selected period, for example, when available memory falls below a threshold. Simmons, however, does not teach sending a *predetermined* number of packets from the transmitting station to the switch, and generating an acknowledge upon receipt of the last packet in the predetermined number of packets in which a time for resuming transmission can be included. In fact, in Simmons, the transmitting station continues sending data frames until a PAUSE signal, if any, is received.

In addition, similar to Stevenson, Simmons is directed to managing flow of data between devices in a network, and not to controlling data flow between cards within a network device and a central process of that network device.

Thus, claim 6 distinguishes patentably over the combined teachings of Stevenson and Simmons. The same arguments apply with equal force to establish that claim 19 is also patentable over the cited art.

In Paragraph 30, the Office Action rejects claim 7 as being obvious over Stevenson in view of what the Examiner characterizes as Applicant's admissions further in view of Simmons and U.S. Patent No. 6,167,029 of Ramarkrishnan.

Claim 7 depends on claim 6, which in turn depends on claim 1. As noted above, Stevenson and Simmons fail to teach the salient features of claim 6. Further, Ramakrishnan deos not cure the shortcomings of Stevenson and Simmons in that it does not teach or suggest transmitting a *predetermined* number of packets from a sender to a received followed by

transmission of an acknowledgement from the receiver to the sender. Rather, similar to Simmons, it discloses utilizing PAUSE frames for providing flow control.

New Claims

Independent method claim 26 recites a method of managing retrieval of statistical data from at least one application executing on a network device that includes linking the application to a usage data monitoring library (UDML) executing on that network device. The UDML periodically polls said application to gather statistical data regarding selected aspects of its performance. The gathered statistical data is transmitted to a usage data server (UDS) executing on that network device in one or more groups of packets, wherein each group has a number of packets less than a predefined threshold, and contains an acknowledge request. This is followed by transmitting an acknowledge packet from the UDS to the UDML upon receipt of the acknowledge request by the UDS to instruct the UDML to transmit a subsequent group of packets.

Support for new claim 26 can be found on pages 76-86 of the application, and throughout the remainder of the specification. Thus, no new matter is added.

None of the cited references teaches or suggests the subject matter of claim 26. In particular, none of the cited references teaches a method for managing retrieval of statistical data from applications executing on a network device, much less specifics of claim 1 and their concomitant advantages. For example, none of the references teaches employing a UDML, linked to an application, to gather and send statistical data associated with that application to a UDS for processing.

CONCLUSION

In view of the above amendments and remarks, Applicants respectfully request reconsideration and allowance of the application. The Examiner is invited to call the undersigned at (617) 439-2514 if there are any remaining issues.

Dated: October 8, 2004

Respectfully submitted,

By _____

Reza Mollaaghababa
Registration No.: 43,810
NUTTER MCCLENNEN & FISH LLP
World Trade Center West
155 Seaport Boulevard
Boston, Massachusetts 02210-2604
(617) 439-2000
(617) 310-9000 (Fax)
Attorney for Applicant

1368750

1368750.1